

PATENT

What is claimed is:

1 1. A switch for a packet switched or frame switched network, said switch having
2 conventional routing circuitry, said network having client devices and storage devices and a
3 storage manager that decides to redirect read and/or write transactions on any basis so as to
4 virtualize said storage devices, said switch comprising:

5 one or more ports which are structured to operate to receive one or more redirection
6 commands from said storage manager or another switch in said network coupled to said
7 storage manager containing old address data and new address data, said ports containing
8 circuitry to store said old address data and new address data in one or more look up
9 tables, said switch containing circuitry functioning to compare packet or frame type
10 information included in a header of at least some packets or frames arriving at a port
11 coupled to a client device or a storage device to determine if the packet or frame is of a
12 type which is to be relabelled and redirected, and, if the packet or frame is of a type to
13 be relabelled and redirected, comparing at least some of the old address data therein to
14 the old address data entries for one or more redirection commands stored in one or more
15 of said look up tables, and if there is a match to any of the old address data of said one or
16 more redirection commands, relabelling said packet or frame by substituting the
17 corresponding new address data from said matching redirection command(s) for said old
18 address data in said packet or frame header, said new address data being such as to cause
19 said relabelled packet or frame to be routed so as to bypass said storage manager but to
20 look like it came from said storage manager, and then passing said relabelled packet or
21 frame to said conventional routing circuitry for forwarding to the destination identified
22 in said new address data, and wherein the sequence in which the type of incoming packet
23 is determined and the old address data lookup is performed is not critical.

1 2. The apparatus of claim 1 wherein each said port contains its own redirection
2 circuitry and its own look up table for storing redirection commands, said redirection
3 circuitry in every port being structured to receive first and second redirection commands
4 for every write transaction to be redirected and to use configuration data defining which
5 switch ports are coupled to which devices, said configuration data being stored in said switch
6 or received in a message from said storage manager and said configuration data being used to

PATENT

7 forward said first redirection command to a first port of said switch coupled to the client
8 device which originated said write request for storage in a look up table, and to forward the
9 second redirection command to a second port coupled to the storage device which is to store
10 the data to be written for storage in a look up table, said redirection circuitry being
11 structured to receive a redirection command for every read transaction to be redirected and
12 to use configuration data to forward said redirection command to a third port coupled to the
13 storage device where the requested data is stored for storage in a look up table.

1 3. The apparatus of claim 1 wherein each said port contains its own redirection
2 circuitry and its own look up table for storing redirection commands, said redirection
3 circuitry in every port being structured to receive first and second redirection commands
4 for every write transaction to be redirected and to forward said first and second redirection
5 commands to all ports of said switch for storage in the look up tables of said ports, said
6 redirection circuitry also being structured to receive a redirection command for every read
7 transaction to be redirected and to forward said redirection command to all ports of said
8 switch for storage in a look up table in said third port.

1 4. The apparatus of claim 1 wherein each said port contains its own redirection
2 circuitry, said redirection circuitry in every port being structured to receive first and
3 second redirection commands for every write transaction to be redirected and to forward
4 said first and second redirection commands to a shared look up table in said switch to which
5 all redirection circuits in all ports have access, said redirection circuitry also being
6 structured to receive a redirection command for every read transaction to be redirected and
7 to forward said redirection command for storage in said shared look up table.

1 5. The apparatus of claim 1 wherein said switch contains shared redirection circuitry
2 that performs the redirection process and a shared look up table for storing redirection
3 command used by said shared redirection circuitry to perform said redirection process, and
4 wherein each said port contains circuitry to store redirection commands in said shared look
5 up table by any method, and to forward frames or packets to said shared redirection
6 circuitry for analysis and relabelling if said frames or packets are to be redirected.

PATENT

1 6. A virtualizing storage manager for a packet switched or frame switched network, said
2 network having client devices and storage devices and a switch having any structure capable
3 of redirecting read and/or write transactions based upon commands from said storage
4 manager so as to virtualize said storage devices, said storage manager comprising:

5 conventional storage manager circuitry to carry out conventional processing of read
6 and/or write transactions that are not redirected; and

7 a computer programmed to or logic circuitry structured to decide whether to
8 redirect packets or frames of a read or write transaction based upon any redirection
9 criterion, and programmed or structured to generate and send at least one said
10 redirection command to a switch in said network for every read or write transaction to
11 be redirected, and programmed or structured to send a corresponding read or write
12 request to the appropriate storage device for every received read or write request from
13 a client device to start the transaction, and programmed or structured to send at least
14 one purge command to a switch in said network which stored said redirection commands
15 in one or more look up tables, said at least one purge command corresponding to one or
16 more completed read or write transactions that have been redirected by said switch to
17 cause said switch to purge from said one or more look up tables the redirection command
18 or commands pertaining to said one or more read or write transaction that has been
19 completed.

1 7. The apparatus of claim 6 wherein said storage manager makes a decision based upon
2 programmable redirection criteria.

1 8. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or
4 transfer ready type which is part of a read or write transaction is to be redirected.

1 9. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or

PATENT

4 transfer ready type which is part of a read or write transaction is to be redirected when the
5 number of outstanding read and/or write requests is greater than or equal to a threshold.

1 10. The apparatus of claim 9 wherein said threshold is programmable.

1 11. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or
4 transfer ready type which is part of a read or write transaction is to be redirected when the
5 number of outstanding read and/or write requests is such that a measured latency between
6 receipt of a read or write request and completion of the corresponding read or write
7 transaction is greater than a threshold latency.

1 12. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or
4 transfer ready type which is part of a read or write transaction is to be redirected when said
5 storage manager decides based upon any type criteria or any type measurement that it has
6 become a bottleneck.

1 13. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or
4 transfer ready type which is part of a read or write transaction is to be redirected when the
5 count of a counter which was started when a read or write request from a client device was
6 received and stopped when the read or write transaction corresponding to said request has
7 been completed exceeds a predetermined threshold.

1 14. The apparatus of claim 13 wherein said predetermined threshold is a number read
2 from a look up table which stores threshold values based upon the size of the read or write
3 transaction, and wherein said means for deciding performs the following steps:

PATENT

(1) analyzing each read or write request to determine the size of the transaction and retrieves the appropriate latency threshold value from said look up table based upon the size of the read or write transaction;

(2) starting a counter when said read or write request was received;

(3) perform conventional read or write transaction processing including sending a read or write command to the appropriate storage device and including forwarding data and transfer ready frames to the client or storage device as appropriate to the particular type of transaction being performed;

(4) stop said counter when said conventional read or write transaction was completed;

(5) comparing the counter value when the transaction was completed to the threshold number looked up from said look up table as a measure of the time it took to complete the conventional read or write transaction; and

(6) if the time taken to complete the conventional read or write transaction exceeded the threshold, generating one or more redirection commands for subsequent read or write requests and sending them to the switch;

(7) performing any method of stopping redirection after some time passes and resuming the latency measuring process of steps (1) through (6) for subsequent read or write requests until a latency threshold value is again exceeded, and then repeating steps (6) and (7).

15. The apparatus of claim 14 wherein step (7) comprises continuing redirection for a fixed or programmable amount of time, and then ceasing redirection and resuming performance of steps (1) through (6) for subsequent read or write requests until a latency threshold value is again exceeded, and then repeating steps (6) and (7).

16. The apparatus of claim 14 wherein step (7) comprises continuing redirection for all subsequent read or write requests except every Xth read or write request where X is any integer and carrying out conventional read or write transaction processing as appropriate for every Xth read or write request and measuring the latency of said Xth request by performance of steps (1) through (6), and if the latency threshold has not been exceeded by said Xth request, resuming performance of steps (1) through (6) for subsequent read or

PATENT

7 write requests until a latency threshold value is again exceeded, and then repeating steps
8 (6) and (7).

1 17. The apparatus of claim 6 wherein said storage manager makes a decision whether or
2 not to redirect data and transfer ready packets or frames of a particular read or write
3 request based upon redirection criteria that is that every packet or frame of a data or
4 transfer ready type which is part of a read or write transaction that exceeds a certain size
5 as determined from the original read or write request is to be redirected.

1 18. The apparatus of claim 6 wherein said storage manager includes a cache memory and
2 a cache algorithm that manages said cache memory, and wherein said storage manager makes
3 a decision whether or not to redirect data and transfer ready packets or frames of a
4 particular read or write request based upon whether the requested data is already stored in
5 said cache memory in said storage manager, or makes a decision whether or not to redirect a
6 read transaction based upon whether said storage manager cache algorithm wants to copy the
7 data of the read request into said cache memory in said storage manager.

1 19. The apparatus of claim 1 wherein each port of said switch is structured to receive
2 redirection commands either from an external source or from a source internal to the
3 switch such as the routing circuitry and store all said redirection commands in a look up
4 table maintained by said port and use said look up table to do all necessary relabelling and
5 redirection operations.

1 20. A switch for a packet switched or frame switched network, said switch having
2 conventional routing circuitry, said network having client devices and storage devices and a
3 storage manager coupled to said switch, said switch comprising:

4 one or more ports which are structured to operate to receive redirection commands
5 from a storage manager containing old address data and new address data and to store the
6 old address data and new address data contained therein in a shared look up table, but
7 otherwise to function as conventional switch ports;

8 shared redirection interface circuitry coupled to each of said ports and said shared
9 look up table and coupled to said conventional routing circuitry, said shared redirection

PATENT

interface circuitry structured to compare packet or frame type information included in a header of each packet or frame arriving from a port coupled to a client device or a storage device to determine if the packet or frame is of a type which is to be relabelled and redirected, and, if the packet or frame is of a type to be relabelled and redirected, then comparing the old address data therein to the old address data in said shared look up table, and if there is a match to any of the old address data entries in said look up table, substituting the corresponding new address data from said look up table, and forwarding said packet or frame to said conventional routing circuitry for routing.

21. A packet switched or frame switched network, comprising:

one or more client devices;

one or more storage devices;

a storage manager having redirection circuitry functioning to receive read or write transaction requests from said one or more client devices and redirect at least some of said read and write transaction requests issued by said client devices by issuing redirection commands which contain old address data and new address data which will cause relabelling of data and transfer ready frames so that they are routed between said client device which issued said request and the storage device which is to be read from or written to so as to bypass said storage manager, said redirection circuitry also for sending a corresponding read or write request to the appropriate storage device for every received read or write request from a client device to start the transaction, said redirection circuitry also functioning to receive status packets or frames indicating particular read or write transactions that have been redirected have been completed and for issuing purge commands which function to cause purging of old address and new address data of redirected read or write transactions that have been completed;

one or more switches, each having conventional routing circuitry and each further comprising:

a shared lookup table;

a shared redirection interface circuit;

one or more ports which are structured to operate to receive redirection commands from said storage manager or another switch and to store the old address data and new address data contained therein in said shared look up table, and

PATENT

structured to receive said purge commands and to use the information therein to purge obsolete old address data and new address data from said shared look up table for read or write transactions that have been completed, but otherwise to function as conventional switch ports;

and wherein said shared redirection interface circuitry coupled to each of said ports and said shared look up table and coupled to said conventional routing circuitry, said shared redirection interface circuitry structured to compare packet or frame type information included in a header of each packet or frame arriving from a port coupled to a client device or a storage device to determine if the packet or frame is of a type which is to be relabelled and redirected, and, if the packet or frame is of a type to be relabelled and redirected, then comparing the old address data therein to the old address data in said shared look up table, and if there is a match to any of the old address data entries in said look up table, substituting the corresponding new address data from said look up table, and forwarding said packet or frame to said conventional routing circuitry for routing; and

a data path coupling each said client devices, storage devices and said storage manager to at least one of said one or more switches and coupling said one or more switches together.

22. A packet switched or frame switched network, comprising:

one or more client devices;

one or more storage devices;

a storage manager having redirection circuitry functioning to receive read or write requests from said one or more client devices and redirect at least some of said read and write transaction requests issued by said client devices by issuing redirection commands which contain old address data and new address data which will cause relabelling of data and transfer ready frames so that they are routed between said client device which issued said request and the storage device which is to be read from or written to so as to bypass said storage manager, said redirection circuitry also for sending a corresponding read or write request to the appropriate storage device for every received read or write request from a client device to start the transaction, said redirection circuitry also functioning to receive status packets or frames indicating particular read or write transactions that

PATENT

14 have been redirected have been completed and for issuing purge commands which
15 function to cause purging of old address and new address data of redirected read or write
16 transactions that have been completed;

17 one or more switches each having conventional routing circuitry and each further
18 comprising one or more ports

19 and wherein at least the ports coupled to said storage manager or another
20 switch being structured to operate to receive one or more redirection commands
21 from said storage manager or another switch in said network coupled to said storage
22 manager containing old address data and new address data and store said old address
23 data and new address data in one or more look up tables and to receive purge
24 commands from said storage manager or another switch in said network and use the
25 information therein to purge old and new address data pertaining to redirected read
26 or write transactions that have been completed from said one or more look up tables,

27 each said switch containing redirection circuitry functioning to compare
28 packet or frame type information included in a header of at least some packets or
29 frames arriving at a port coupled to a client device or a storage device to determine if
30 the packet or frame is of a type which is to be relabelled and redirected, and, if the
31 packet or frame is of a type to be relabelled and redirected, then comparing the old
32 address data therein to the old address data entries in said look up table, and if there
33 is a match to any of the old address data entries in said one or more look up tables,
34 substituting the corresponding new address data from said look up table for said old
35 address data in said packet or frame header and passing said packet or frame to said
36 conventional routing circuitry for forwarding to the destination identified in said
37 new address data, at least the port; and

38 a data path coupling each said client devices, storage devices and said storage manager
39 to at least one of said one or more switches and coupling said one or more switches
40 together.

1 23. The apparatus of claim 21 or 22 wherein said redirection circuitry in said storage
2 manager includes means to decide which read and/or write transactions to redirect based
3 upon predetermined redirection criteria which could be a single criteria or any one of a

PATENT

4 plurality of different criteria, said criteria being either fixed or programmable, and which
5 particular criteria is used can be fixed or configurable.

1 24. The apparatus of claim 22 wherein said one or more look up tables comprise a look
2 up table in each port and wherein said redirection circuitry comprises a redirection circuit
3 in every port.

1 25. The apparatus of claim 22 wherein said one or more look up tables comprise a look
2 up table in each port and wherein said redirection circuitry comprises a redirection circuit
3 in every port, and wherein said redirection circuitry in said storage manager includes
4 means to decide which read and/or write transactions to redirect based upon predetermined
5 redirection criteria which could be a single criteria or any one of a plurality of different
6 criteria, said criteria being either fixed or programmable, and which particular criteria is
7 used can be fixed or configurable.

1 26. A switch for a packet or frame switched network including one or more client
2 devices and one or more storage devices and a storage manager server, said switch having
3 conventional routing circuitry and means for receiving and storing redirection commands
4 from said storage manager server and for comparing address data of predetermined types of
5 frames or packets arriving at said port to stored redirection data from said redirection
6 commands and for readdressing any frames or packets of said predetermined types which
7 match any stored redirection command and for receiving purge commands from said storage
8 manager server and purging old address data and new address data identified in said purge
9 commands of redirection commands issued for transactions that have been completed.

1 27. A storage manager for a packet or frame switched network including one or more
2 client devices and one or more storage devices and a switch, comprising:
3 means receiving read or write transaction requests from said one or more client
4 devices and mapping the data designated in said read or write request to a storage device,
5 and for sending corresponding read or write requests to said storage device to start the
6 transaction; and

PATENT

7 means for using redirection criteria to decide which read and write requests from
8 said client devices to redirect and for generating and sending redirection commands to
9 said switch for the read and write requests to be redirected, said redirection commands
10 containing old address data and new address data which will cause data and transfer ready
11 frames of read and write transactions to be redirected to be relabelled so that they are
12 routed by said switch to bypass said storage manager but look like they came from said
13 storage manager; and

14 means for receiving status frames or packets and for issuing purge commands to said
15 switch when said status frames or packets indicate a read or write transaction which has
16 been redirected has been completed, said purge commands for causing said switch to
17 purge from one or more look up tables in said switch old address data and new address
18 data identified in said purge commands of redirection commands previously issued for
19 ~~transactions that have been completed.~~

1 28. A packet or frame switched network including one or more client devices and one or
2 more storage devices each of which is coupled to a packet or frame switch, comprising:

3 a switch having conventional routing circuitry therein, said switch
4 comprising:

5 one or more ports in said switch, said switch including redirection
6 circuitry structured to operate to receive redirection commands from a
7 storage manager and to store the redirection address data containing old
8 address data and new address data contained therein in a look up table, said
9 redirection circuitry also structured to receive incoming packets or frames
10 and determine if they are data or transfer ready type frames or packets and,
11 if not to transfer said packet or frame to said conventional routing circuitry
12 but if so, to compare the address data therein to the old address data in said
13 look up table, and, if there is a match to any of the old address data in said
14 look up table, functioning to substitute the new address data from said look up
15 table from the entry which includes old address data that matched the address
16 data in the incoming packet or frame and pass the packet or frame to said
17 conventional routing circuitry; and

PATENT

1 8 a storage manager comprising a one or more port adapter circuits, a memory and a
1 9 microprocessor coupled to said memory and programmed to receive read or write
2 0 requests from said one or more client devices and map the data designated in said read or
2 1 write request to a particular storage device and to send corresponding read or write
2 2 requests to said storage device to start the transaction, and programmed to determine
2 3 from the frame or packet header information of said read or write request the specific
2 4 client device from which the request came and how large the transaction is, and, if the
2 5 transaction is larger than a redirection size criteria, for generating and sending one or
2 6 more redirection commands to said switch for storage in at least one look up table
2 7 maintained therein, said redirection command containing old address data and new
2 8 address data that will cause said switch to receive data and transfer ready frames for
2 9 transactions to be redirected to relabel said packet or frame with new address data so
3 0 that said data or transfer ready frames will be routed by said switch so as to bypass said
3 1 storage manager but look like they came from said storage manager, and for performing
3 2 conventional processing on any read or write request that is not to be redirected, and
3 3 programmed or structured to monitor status frames to determine when a redirected
3 4 transaction has been completed and sending purge commands to said switch to purge the
3 5 old and new address data from said look up table for read and write transactions that have
3 6 been completed; and
3 7 a data path coupling said switch to said storage manager and to said client and storage
3 8 devices.

1 29. The apparatus of claim 28 wherein each said port maintains its own redirection
2 circuitry and its own look up table storing redirection commands, and wherein said
3 microprocessor in said storage manager server looks up the specific port to which a specific
4 storage device involved in a transaction is coupled and/or the specific port to which the
5 client device involved in a specific transaction is coupled and sends said one or more
6 redirection commands for the transaction only to the appropriate specific port.

1 30. The apparatus of claim 28 wherein each said port maintains its own look up table
2 storing redirection commands, and wherein said microprocessor in said storage manager
3 server sends all said redirection commands to all ports in said switch and wherein each port

PATENT

4 includes redirection circuitry to store redirection commands in the port's look up table and
5 use the data therein to relabel for redirection data and transfer ready frames.

1 31. The apparatus of claim 28 wherein said switch maintains a shared look up table
2 storing all redirection commands received from said storage manager, said shared look up
3 table for use by all ports, and wherein said microprocessor in said storage manager sends
4 all said redirection commands to a port to which said storage manager server is coupled with
5 instructions to forward said redirection command for storage in said shared look up table,
6 and wherein each said port of said switch has circuitry to use said shared look up table to
7 determine if an incoming frame or packet needs to be relabelled for redirection and at least
8 the ports coupled to said storage manager include circuitry to store all said redirection
9 commands in said shared look up table.

1 32. The apparatus of claim 28 wherein said microprocessor in said storage manager
2 server is programmed to sum the number of outstanding read requests that have not yet been
3 fulfilled, and only issues redirection commands when the number of outstanding read
4 requests exceeds a threshold which is either fixed or programmable and said size is greater
5 than said redirection size criteria.

1 33. The apparatus of claim 28 wherein said microprocessor in said storage manager
2 server is programmed to monitor the time between receipt of a read or write request and
3 completion thereof, and only issues redirection commands when the time to fulfillment
4 exceeds a threshold which is either fixed or programmable and the size of said transaction is
5 greater than said redirection size criteria.

1 34. A packet or frame switched network including one or more client devices and one or
2 more storage devices, each of which is coupled to a packet or frame switch, comprising:
3 a switch having conventional routing circuitry therein and further comprising:
4 shared redirection circuitry;
5 a shared look up table;
6 one or more ports in said switch which are structured to operate to
7 receive relabelling and redirection and purge commands from a storage

PATENT

8 manager and to forward them to said shared redirection circuitry for storage
9 in said shared look up table and structured to forward data and transfer ready
10 frames to said shared redirection circuitry and to forward all other types of
11 packets or frames to said conventional routing circuitry;

12 and wherein said shared redirection circuit includes circuitry to
13 receive said redirection commands and store them in said shared look up table
14 and to receive incoming data or transfer ready packets or frames from said
15 ports and compare the address data therein to the old address data in said look
16 up table, and, if there is a match to any of the old address data in said look up
17 table functioning to substitute the new address data from the matching entry
18 in said look up table for the old address data in the incoming packet or frame
19 and pass the packet or frame to said conventional routing circuitry, said
20 shared redirection circuit also structured to receive said purge commands
21 and use the information therein to purge from said shared look up table the
22 old and new address data of redirection commands for completed read or write
23 transactions; and

24 a storage manager including conventional cache memory and circuitry programmed or
25 structured to receive read or write requests from said one or more client devices identifying
26 specific data to be read and to map the identified data to a storage device and, for each read or
27 write request, send a corresponding read or write requests to said storage device to get the
28 transaction started, and for storing the most frequently requested data blocks from read
29 transactions in said cache memory and programmed or structured to monitor read requests
30 for read requests for blocks of data that are stored in said cache memory, and, if a requested
31 block of data is stored in said cache memory, for sending the requested data from the cache
32 memory back to the client device which requested the data instead of sending a read request
33 to a storage device that stores the requested data, and further comprises circuitry to decide
34 whether to redirect a particular read or write request based upon fixed or programmable
35 redirection criteria, and programmed or structured to determine from the frame or packet
36 header information of said packet or frame the specific client device from which the request
37 came, and, if said read or write request is not to be redirected, for performing conventional
38 processing to complete said read or write request, but if said read or write request is to be
39 redirected, for generating and sending one or more redirection commands to said switch for

PATENT

40 storage in said shared look up table maintained therein, said redirection command containing
41 old address data and new address data that will cause said shared redirection circuitry that
42 receives a data or transfer ready packet or frame that is part of the transaction to be
43 redirected to relabel said packet or frame with new address data and forward said packet or
44 frame to said conventional routing circuitry such that said packet or frame gets routed so as
45 to bypass said storage manager but so as to look like it came from said storage manager, and
46 programmed or structured to monitor status frames to determine when a redirected
47 transaction has been completed and sending purge commands to said switch to purge the old
48 and new address data from said shared look up table for read and write transactions that have
49 been completed; and

50 a data path coupling said switch to said client and storage devices and to said storage
51 manager.

1 35. A packet or frame switched network including one or more client devices and one or
2 more storage devices, comprising:

3 a data path coupled to said client devices and said storage devices;

4 a storage manager coupled to said data path;

5 a switch having conventional routing circuitry and including:

6 one or more ports in said switch which contain redirection circuitry to
7 receive relabelling and redirection commands from a storage manager and to store
8 the redirection address data containing old address data and new address data
9 contained therein in at least one look up table and to receive incoming data and
10 transfer ready packets or frames and comparing the address data therein to the old
11 address data in said look up table and to compare packet or frame type information
12 included in a header of said packet or frame, and, if there is a match to any of the old
13 address data in said look up table and the type information matches the types of
14 packets or frames that are to be redirected, functioning to substitute at least the new
15 address data from said look up table from the entry which includes old address data
16 that matched the address data in the incoming packet or frame and pass the packet or
17 frame to said conventional routing circuitry and performing conventional processing
18 to route all other packets or frames to said conventional routing circuitry; and

19 wherein said storage manger includes:

PATENT

20 means for receiving write and read requests from any of said one or more client
21 devices identifying specific data to be read or written and for mapping said read and
22 write requests to particular storage devices, and for sending corresponding read or
23 write requests to said storage devices to get the transaction started, and for determining
24 if the transaction is to be redirected, and, if not, for performing conventional processing
25 to complete the transaction, but if said transaction is to be redirected, for analyzing said
26 read and write requests to determine how many blocks of data are to be read or written,
27 and, if the number of blocks is greater than a threshold, which can be fixed or
28 programmable, for mapping the requested blocks to a storage device which stores the
29 requested data and for generating and sending one or more redirection commands to said
30 switch for storage in said at least one look up table maintained therein, said redirection
31 commands containing old address data and new address data that will cause the redirection
32 circuitry in the port that receives a data or transfer ready packet or frame to relabel
33 said packet or frame with new address data that will cause said packet or frame to be
34 routed by said conventional routing circuitry so as to bypass said storage manager but
35 look like it came from said storage manager, and for monitoring status packets or frames
36 to determine when a redirected transaction has been completed and sending purge
37 commands to said switch to purge the old and new address data from said at least one look
38 up table for read and write transactions that have been completed.

1 36. A process of redirecting data frames or packets in a packet or frame switched
2 network comprising the steps:
3 receiving at a storage manager from a client computer one or more read requests;
4 receiving from a storage manager at a packet or frame switch at least one redirection
5 command to redirect the data frames or packets of at least some of said read requests
6 originated by a client device and storing said redirection commands;
7 transmitting corresponding read requests from said storage manager to one or more
8 storage devices;
9 responding to said read requests by transmitting frames or packet containing the
10 requested data addressed to said storage manager; and
11 in said switch, intercepting said frames or packets of data responding to read
12 requests, and, if a redirection command is stored for a read transaction which certain

PATENT

13 packets or frames containing the requested data are a part of, changing the header of the
14 frames or packets of said particular read transaction which contain the requested read
15 data so as to reroute them through said switch so as to bypass said storage manager but
16 relabeling the headers thereof such that the rerouted packets or frames appear to
17 originate from said storage manager as a response to said read request.

1 37. A process of redirecting data frames or packets in a packet or frame switched
2 network, said process carried out in a switch and comprising the steps:
3 receiving from a storage manager at a packet or frame switch redirection commands
4 to redirect the data frames or packets which are responses to at least some of said read
5 requests originated by a client device and storing said redirection commands; and
6 in said switch, intercepting frames or packets of read data responding to said at least
7 some read requests, and, if a redirection command is stored for a read transaction which
8 certain packets or frames containing the requested read data are a part of, changing the
9 header of the frames or packets of said particular read transaction which contain the
10 requested read data so as to reroute them through said switch so as to bypass said storage
11 manager but relabeling the headers thereof such that the rerouted packets or frames
12 appear to originate from said storage manager as a response to said read request.

1 38. A process for controlling redirection of data frames or packets in a packet or frame
2 switched network, said process carried out in a storage manager and comprising the steps:
3 receiving at a storage manager from a client computer one or more read requests;
4 composing and transmitting to a packet or frame switch redirection commands to
5 redirect the data frames or packets of at least some of said read requests originated by a
6 client device such that said switch relabels data frames for routing so as to bypass said
7 storage manager in travelling to the client device which made the read request but make
8 said relabelled data frame look like it came from said storage manager;
9 transmitting said read requests from said storage manager to one or more storage
10 devices.

1 39. A process of redirecting data frames or packets in a packet or frame switched
2 network, said process carried out in a switch and comprising the steps:

PATENT

3 receiving from a storage manager at a packet or frame switch redirection commands
4 to redirect the data frames or packets which are responses to at least some of said read
5 requests originated by a client device and storing said redirection commands; and

6 in said switch, intercepting frames or packets of read data responding to said at least
7 some read requests, and, if a redirection command is stored for a read transaction which
8 certain packets or frames containing the requested read data are a part of, changing the
9 header of the frames or packets of said particular read transaction which contain the
10 requested read data so as to reroute them through said switch so as to bypass said storage
11 manager but relabeling the headers thereof such that the rerouted packets or frames
12 have the Port_ID of the client that made the read request as the destination address and
13 the Port_ID of the storage device as the source address and have the originator exchange
14 ID originally assigned by said client; and

15 in said client, receiving said frames and mapping them to the original read requests
16 sent to said storage manager.

17 40. A process of redirecting data frames or packets in a packet or frame switched
network comprising the steps:

3 receiving at a storage manager from a client computer a write request;

4 receiving from a storage manager at a packet or frame switch at least two redirection
5 commands to redirect the data and transfer ready frames or packets of said write request
6 and storing said redirection commands;

7 transmitting a corresponding write request from said storage manager to a storage
8 device to get the transaction started;

9 at said storage device, responding to said write request by transmitting one or more
10 transfer ready frames from said storage device when it is ready to receive at least some
11 of the write data, said transfer ready frame(s) or packet(s) addressed to said storage
12 manager; and

13 in said switch, intercepting said transfer ready frame(s) or packet(s) of data
14 responding to write request, and, if a redirection command is stored for a write
15 transaction which said transfer ready frame(s) or packet(s) are a part of, changing the
16 header of the transfer ready frame(s) or packet(s) of said write transaction to be
17 redirected and routing transfer ready frame(s) or packet(s) so as to reroute them

PATENT

18 through said switch so as to bypass said storage manager but so as to appear that transfer
19 ready frame(s) or packet(s) originated from said storage manager;

20 at said client device, receiving transfer ready frame(s) or packet(s) and responding
21 by sending one or more frames or packets of write data to said switch;

22 at said switch, receiving said one or more frames or packets of write data and if a
23 redirection command is stored for a write transaction which said write data frame(s) or
24 packet(s) are a part of, changing the header of the write data frame(s) or packet(s) of
25 said write transaction to be redirected and routing said write data frame(s) or packet(s)
26 so as to bypass said storage manager but so as to appear that write data frame(s) or
27 packet(s) originated from said storage manager; and

28 continuing the process re-routing said transfer ready and write data frames until
29 the entire write transaction is completed.

41. A process of redirecting data frames or packets in a packet or frame switched
2 network, said process carried out in a switch and comprising the steps:

3 receiving from a storage manager at a packet or frame switch one or more
4 redirection commands to redirect the data and transfer ready frames or packets which
5 are part of a write transaction originated by a client device and storing said redirection
6 commands; and

7 in said switch, intercepting write data and transfer ready frames or packets which
8 are part of a write transaction to be redirected and, if one or more redirection commands
9 have been stored for the write transaction of which said write data and transfer ready
10 frames or packets are a part, changing the header of the write data and transfer ready
11 frames or packets so as to reroute them through said switch so as to bypass said storage
12 manager but so as to look like they came from said storage manager and so as to have the
13 originator exchange ID originally assigned by said client to the write transaction; and

14 transmitting said write data and transfer ready frames or packets to said client or
15 said storage device as appropriate.

42. A system comprising:

1 a data path or paths;

2 one or more client devices coupled to said data path(s);
3

PATENT

4 one or more storage devices coupled to said data path(s);

5 one or more switch means coupled to said data path(s) for receiving packets or

6 frames and routing them appropriately according to header address information, and for

7 receiving redirection commands and storing them in one or more lookup tables and for

8 using address data in said redirection commands to redirect predetermined packets or

9 frames; and for receiving purge commands and for using the data in said purge command

10 to purge from said one or more lookup tables said redirection commands for read and/or

11 write transactions that have been completed ; and

12 storage manager means for receiving read and/or write requests from said one or

13 more client devices and for determining by any means whether or not to redirect said

14 read and/or write requests, and, if a read and/or write request is to be redirected, for

15 generating and sending to said switch one or more redirection commands to cause

16 predetermined packets or frames of said read and/or write transaction to be redirected,

17 and for determining when a read and/or write transaction that has been redirected has

18 been completed and for generating one or more purge commands pertaining to said read

19 and/or write transactions that have been completed, and sending said purge commands to

20 said switch to cause it to purge the redirection commands of completed transactions, and,

21 if a read or write transaction is not to be redirected, for processing it conventionally.

1 43. A switch for a packet switched or frame switched network, said switch having

2 conventional routing circuitry, said network having client devices and storage devices and a

3 storage manager that decides to redirect read and/or write transactions on any basis so as to

4 virtualize said storage devices, said switch comprising:

5 one or more ports which are structured to operate to receive one or more redirection

6 commands from said storage manager or another switch in said network coupled to said

7 storage manager containing old address data and new address data, said ports containing

8 circuitry to store said old address data and new address data of said one or more

9 redirection commands in one or more look up tables, said switch containing circuitry

10 functioning to compare packet or frame type information included in a header of at least

11 some packets or frames arriving at a port coupled to a client device or a storage device to

12 determine if the packet or frame is of a type which is to be relabelled and redirected,

13 and, if the packet or frame is of a type to be relabelled and redirected, comparing at least

PATENT

14 some of the old address data therein to the old address data entries for one or more
15 redirection commands stored in one or more of said look up tables, and if there is a
16 match to any of the old address data of said one or more redirection commands,
17 relabelling said packet or frame by substituting the corresponding new address data
18 from said matching redirection command(s) for said old address data in said packet or
19 frame header, said new address data being such as to cause said relabelled packet or
20 frame to be routed so as to bypass said storage manager but to look like it came from said
21 storage manager, and then passing said relabelled packet or frame to said conventional
22 routing circuitry for forwarding to the destination identified in said new address data,
23 each port further comprising circuitry to recognize incoming status frames indicating
24 that a particular read or write transaction has been completed and to respond thereto by
25 automatically purging from one or more lookup tables one or more redirection
26 commands that pertain to the completed transaction of which said status frame is a part
27 and route said status frame to a port of said switch coupled to a storage manager, and
28 wherein the sequence in which the type of incoming packet is determined and the old
29 address data lookup is performed is not critical.